



Virtual Leverage: Server Consolidation in Open Source Environments

Margaret Lewis


Commercial Software Strategist

AMD

What Is Virtualization?



Abstraction of Hardware Components



- Hides complexity of hardware infrastructure, simplifying its management
- Isolates software from underlying hardware to increase reliability and flexibility
- Abstracts services from hardware dependencies for better control and access
- Reduces administration through enhanced automation capabilities

Virtualization technology enables the consolidation of resources

- ✓ **Reduced hardware cost**
 - Higher physical resource utilization
 - Smaller footprint that requires less infrastructure support (power, cooling, space, etc)

- ✓ **Improved flexibility and responsiveness**
 - Resources can be adjusted dynamically
 - Enables *On Demand* and *Adaptive Enterprise* operating environments


Server Virtualization Trends

Consolidating Resources




Carve a Server into Many Virtual Machines

- **Problem:** hardware is underutilized
- **Solution:** partition server into virtual machines that run different applications concurrently
- **Benefit:** more efficient use of hardware resources through workload consolidation




Unite Many Servers into a Virtual Machine

- **Problem:** applications need large capacity
- **Solution:** consolidate commodity computers into a virtual machine that can be reconfigured as needed to run required applications
- **Benefit:** flexibility to resize hardware resource to fit the demands





Unite Many Servers into a Virtual Machine




- Virtualization software manages resources of many servers
- Unites these resources to appear as one server to guest OS and applications
- Might require specific type of server and/or interconnection technology
- Egenera PAN Manager

Carve a Server into Many Virtual Machines




- Available on mainframes and high-end Unix servers for many years
- Applications run independently from each other in separate partitions
- Hardware resources are divvied up according to the needs of the applications running on the system
- Example: Solaris 10 Containers

Carve a Server into Many Virtual Machines



- Virtualization software manages resources between Host/Guest OS's
- Application can suffer decreased performance due to added overhead
- Example: VMware Workstation, VMware GSX



- Hypervisor is host environment
- Enables better software performance by eliminating some of associated overhead
- If Hardware is available, the Hypervisor can be designed to take advantage of it
- Example: VMware ESX, Xen Open Source Hypervisor




- Virtual machine monitor for x86-based systems
- Paravirtualization approach which requires OS's be ported
 - Linux 2.4 and 2.6 and NETBSD run over Xen
- Applications run unmodified
 - Supports demanding applications like MySQL, Apache and PostgreSQL
- Released under the GNU General Public License

<http://www.cl.cam.ac.uk/Research/SRG/netos/xen/>

- OS must be ported to run over Xen
 - More efficient than schemes that rely on trapping faulting instructions or emulating privileged operating system code
- OS's running over Xen execute in x86 privilege ring 1
 - Prevents guest OS's from using the normal privileged instructions to turn on/off interrupts, change page table bases etc.
 - Guest OS make a 'hypercall' down into Xen to ask operations to be performed
- Xen exports specially designed block device and network interface abstractions to guest OS's
 - Results in excellent guest I/O performance

Relative performance on native Linux (L), Xen/Linux (X), VMware Workstation 3.2 (V), and User Mode Linux (U).



<http://www.cl.cam.ac.uk/Research/SRG/netos/xen/performance.html>

AMD64 Processors with Direct Connect Architecture

Integrated Memory Controller

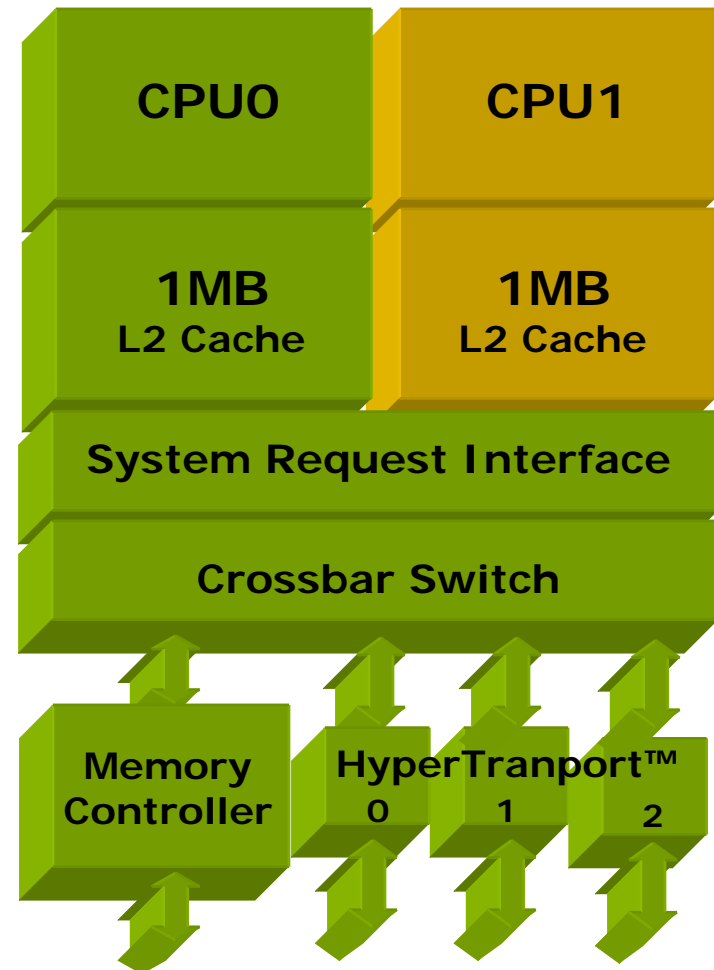
- Increases application performance by reducing memory latency

AMD64 Cores

- Enables both 32- and 64-bit computing
- Eliminates 4GB memory barrier of 32-bit only systems

HyperTransport™ Technology


- Provides up to 24.0 GB's peak bandwidth per processor - reducing I/O bottlenecks
- Directly connects CPUs enabling scalability



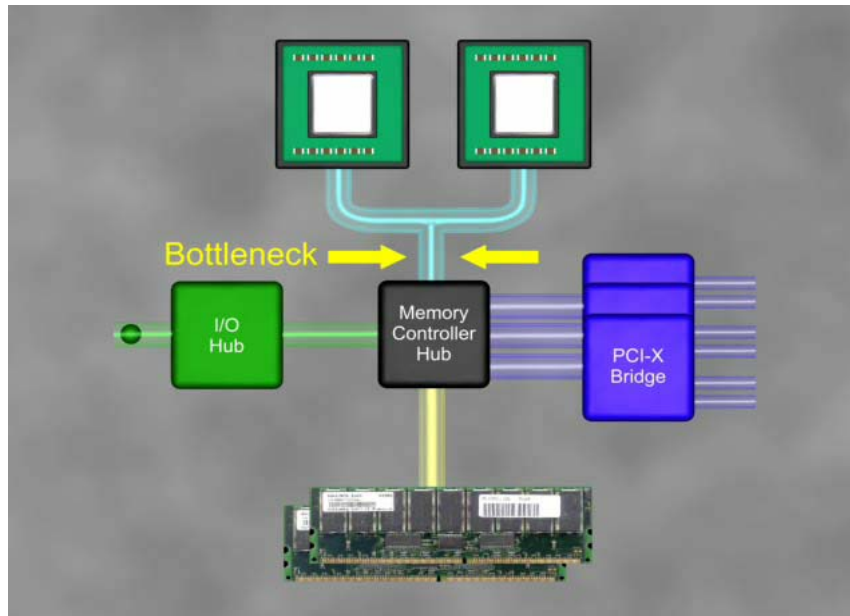
AMD64 Dual-Core Overview

Dual-Core AMD Opteron™ and AMD Athlon™ 64 processors:


- AMD Opteron™ processor **compatible** with existing 940-pin sockets that support 90nm (95W/80A) with a BIOS update, streamlining upgrade paths while increasing performance and value
- AMD Athlon™ 64 processors **compatible** with existing 939-pin sockets with BIOS update
- Completely **compatible** with x86 and AMD64 applications while benefiting multi-threaded environments
- AMD64 technology designed from the ground up for multi-core



Traditional x86 Architecture



Direct Connect Architecture




Direct Connect Architecture revolutionizes the system architecture by eliminating the bottlenecks of the front-side bus

AMD64 processors with Direct Connect Architecture are well suited to support virtualization technology


- ✓ Integrated Memory Controller
 - High-speed, low latency access to memory
 - Enables Host and Guest OS's to run more efficiently
- ✓ HyperTransport™ Technology
 - Improved scalability and I/O capabilities
 - Supports more Guest OS sessions and/or users sessions
- ✓ AMD64 Core
 - Retains compatible with x86 software
 - Provides support for legacy OS's and applications
- ✓ Multi-Core Technology
 - Single physical processor with multiple execution cores
 - Delivers high-bandwidth communication between virtual machines

Web Bench: 4P AMD Opteron™ Processors vs. 4P Xeon 4 VMs per machine (1 VM per CPU)



<http://www.veritest.com/clients/reports/amd/default.asp>

Web Bench: 4P AMD Opteron™ Processors vs. 4P Xeon 8 VMs per machine (2 VM per CPU)



<http://www.veritest.com/clients/reports/amd/default.asp>

AMD64 Processors and Virtualization

Current Software Status



Software	Company	URL	32-bit Linux	64-bit Linux	32-bit Solaris	64-bit Solaris	32-bit Windows®	64-bit Windows
Solaris Containers	Sun Microsystems	http://www.sun.com/2004-0330/feature/index.html			✓	✓		
Trigence Application Environment (AE)	Trigence	http://www.nwfusion.com/newsletters/servers/2004/1206server2.html			✓	✓		
Virtual Iron VFe	Virtual Iron	http://www.virtualiron.com/	✓					
Virtual PC 2004	Microsoft	http://www.microsoft.com/windows/virtualpc/default.mspx					✓	
Virtual Server 2005	Microsoft	http://www.microsoft.com/windowsserver2003/evaluation/trial/virtualserver.mspx					✓	
Virtuozzo	SWsoft	http://www.sw-soft.com/virtuozzo/?ad=google	✓	Beta	Planned	Planned	Beta	Planned
VMware Ace	VMware	http://www.vmware.com/products/desktop/ace_features.html	✓		✓		✓	
VMware GSX Server	VMware	http://www.vmware.com/products/	✓	✓			✓	✓
VMware ESX Server	VMware	http://www.vmware.com/products/	✓	Planned			✓	Planned
VMware SMP	VMware	http://www.vmware.com/products/server/vsmp_features.html	✓	Planned			✓	Planned
VMware Virtual Center	VMware	http://www.vmware.com/products/vmanage/vc_features.html#vmotion	✓				✓	
VMware Workstation	VMware	http://www.vmware.com/products/	✓	✓	✓		✓	✓
Xen Hpvvisor	Open Source	http://www.cl.cam.ac.uk/Research/SRG/netos/xen/	✓	2Q05				

Future Trends in Virtualization



- Native virtualization of x86 architecture requires “unnatural acts” to achieve – leading to increased performance overhead, lower security, and increased complexity.
- Moving functionality traditionally served by software-based hypervisor into the processor helps to solve these problems.
- **“PACIFICA” is AMD’s technology to provide silicon enhanced virtualization.**
- “PACIFICA” allows the software vendors to focus on the value-add, leaving the worry of proper emulation to the processor.

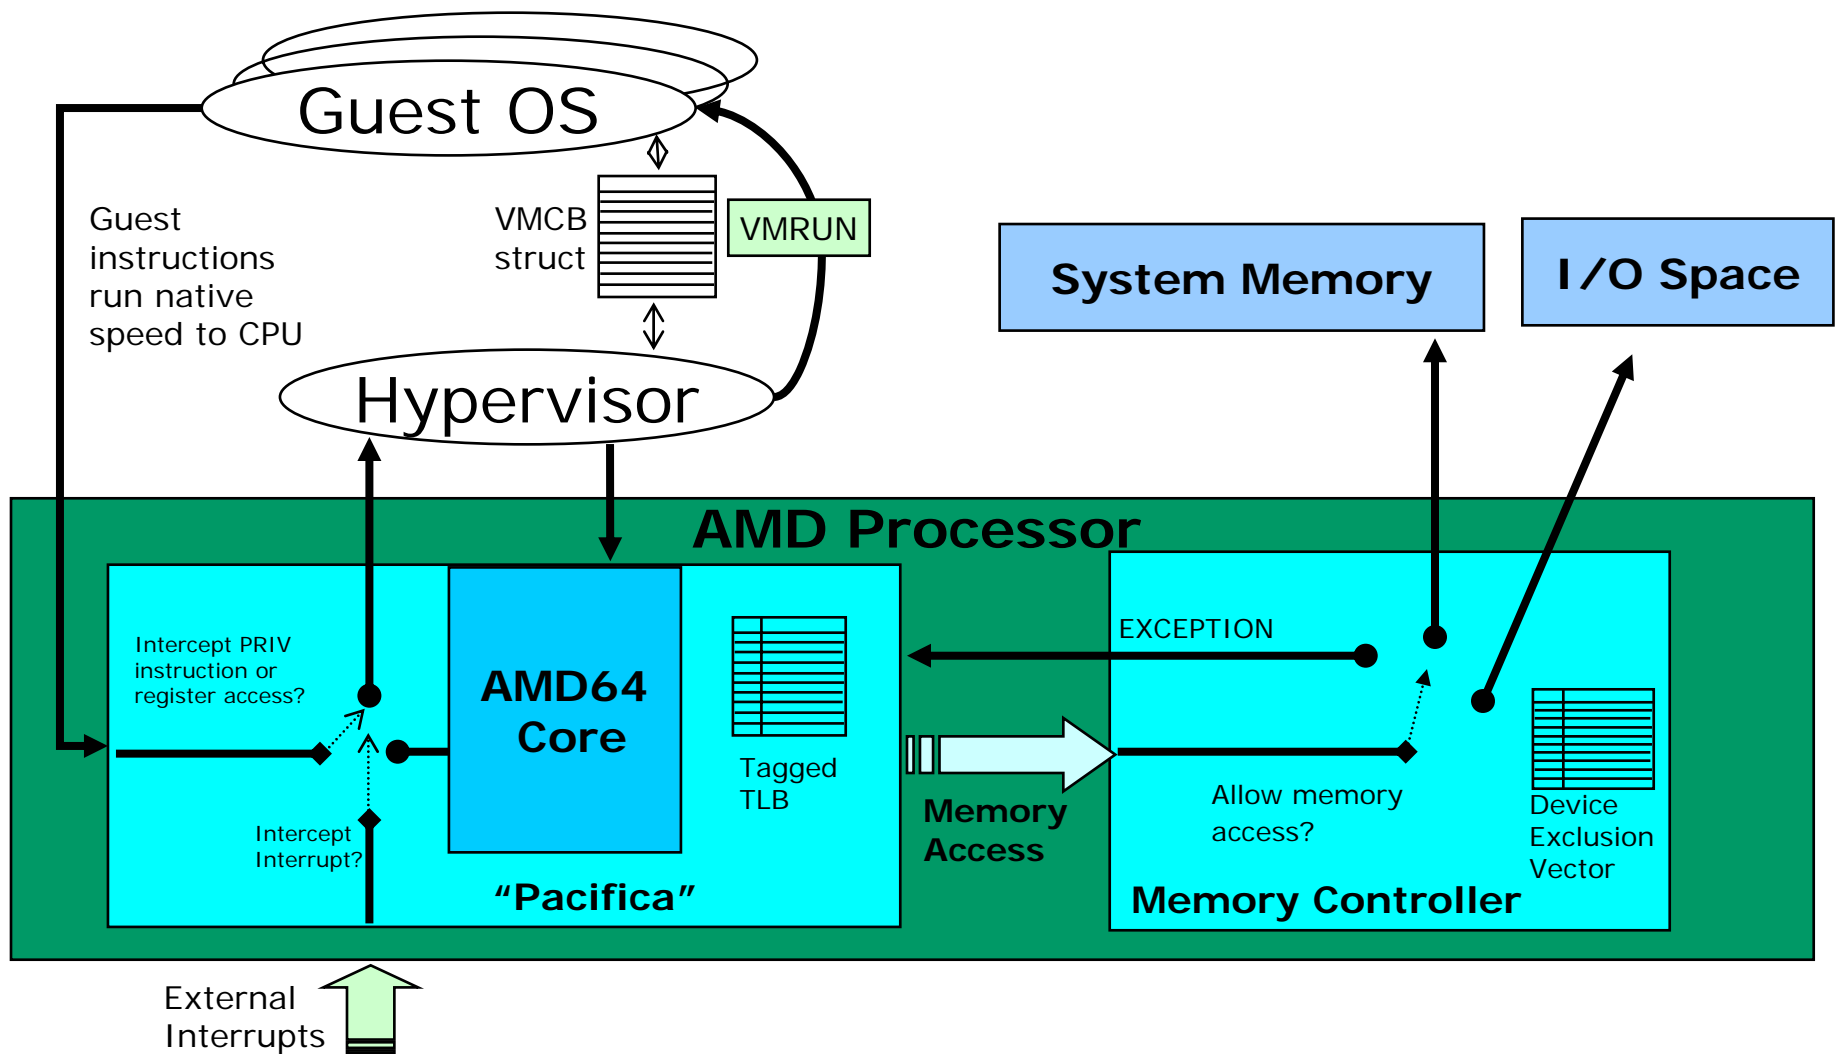
“Pacifica” virtualization technology allows AMD to continue to offer a competitive performance roadmap while meeting the system architecture demands of our customers.

“Pacifica” Overview & Highlights

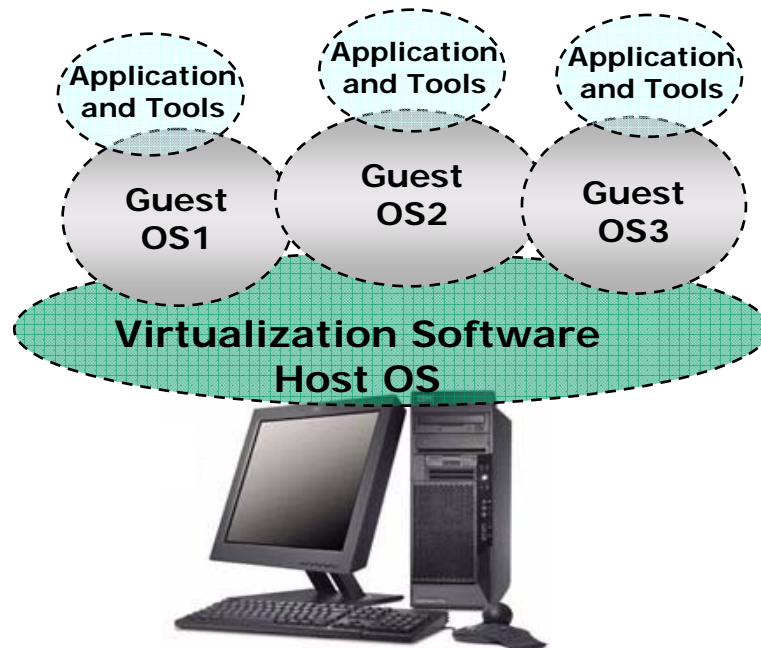


- “**Pacifica**” is the code name for AMD silicon feature set which enhances the performance of native virtualization software
- “**Pacifica**” enabled parts will launch in *all* AMD processors beginning in 2006
- Completely compatible with x86 and AMD64 applications while benefiting **virtualization environments**, where Hypervisor is “Pacifica” aware
- Virtualization and partitioned applications will experience the greatest **performance advantage**. Examples of these applications include:
 - Server consolidation
 - OS migration
 - Blade computing
 - Desktop Hypervisor-based security
- AMD64 processors with *enhanced virtualization* is a continuing example of how AMD is extending it's **Direct Connect Architecture** and **multi-core technology** leadership

"Pacifica" Silicon Enhanced Virtualization



Reduces software development and testing cycles




- Shortening the time needed to provision, install, build, test and restore a new machine
- Development teams can share development environments and pre-packaged operating system and application testing configurations

Summary

Virtual Machine Benefits



Can use more reliable, highly scalable servers to minimize system failures




- Partitions and isolates a server into secure and transportable virtual machines
- Errors that interfere with operation running in one virtual machine have no effect on other virtual machines on the same system
- Obsolete hardware can be upgraded without software losing compatibility

Summary

Virtual Machine Benefits



Allows for more efficient utilization of resources through server consolidation



"If mature virtualization technologies could be applied to x86 servers today, a conservative rough estimate is that overall IT spending in support of x86 servers would decline by 20% to 30%."

- Consolidate multiple applications onto newer server regardless of operating environment
- Consolidate multiple servers to support demanding application
- Helps to reduces costs associated with space, power, installation, integration, and administration

Gartner, Predicts 2004: Server Virtualization Evolves Rapidly

Trademark Attribution



© 2005 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow Logo, AMD Athlon, AMD Opteron, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. HyperTransport is a licensed trademark of the HyperTransport Technology Consortium. Microsoft and Windows are registered trademarks of Microsoft Corporation in the U.S. and/or other jurisdictions. PCI-X is a registered trademark of PCI-SIG. Other names are for informational purposes only and may be trademarks of their respective owners.